ROAR: A QoS-Oriented Modeling Framework for Automated Cloud Resource Allocation and Optimization

Yu Sun[a], Jules White[a], Sean Eade[b], Douglas C. Schmidt[a]

[a]*Department of Electrical Engineering and Computer Science, Vanderbilt University*
[b]*Siemens Industry*

Abstract

Cloud computing offers a fast, easy and cost-effective way to configure and allocate computing resources for web applications, such as consoles for smart grid applications, medical records systems, and security management platforms. Although a diverse collection of cloud resources (e.g., servers) is available, choosing the most optimized and cost-effective set of cloud resources for a given web application and set of quality of service (QoS) goals is not a straightforward task. Optimizing cloud resource allocation is a critical task for offering web applications using a software as a service model in the cloud, where minimizing operational cost while ensuring QoS goals are met is critical to meeting customer demands and maximizing profit. Manual load testing with different sets of cloud resources, followed by comparison of test results to QoS goals is tedious and inaccurate due to the limitations of the load testing tools, challenges characterizing resource utilization, significant manual test orchestration effort, and challenges identifying resource bottlenecks.

This paper introduces our work using a modeling framework - ROAR (Resource Optimization, Allocation and Recommendation System) to simplify, optimize, and automate cloud resource allocation decisions to meet QoS goals for web applications, including complex multi-tier application distributed in different server groups. ROAR uses a domain-specific language to describe the configuration of the web application, the APIs to benchmark

Email addresses: yusun@isis.vanderbilt.edu (Yu Sun), jules@dre.vanderbilt.edu (Jules White), sean.eade@siemens.com (Sean Eade), schmidt@dre.vanderbilt.edu (Douglas C. Schmidt)

and the expected QoS requirements (e.g., throughput and latency), and the resource optimization engine uses model-based analysis and code generation to automatically deploy and load test the application in multiple resource configurations in order to derive a cost-optimal resource configuration that meets the QoS goals.

Keywords: Cloud Computing, Resource Allocation, Resource Optimization, Domain-Specific Language, Load Testing and Benchmarking

---

Type	ECU	Memory (GB)	Price
m3.medium	3	3.75	$0.07 / hour
m3.large	6.5	7.5	$0.14 / hour
m3.xlarge	13	15	$0.28 / hour
m3.2xlarge	26	30	$0.56 / hour
c3.4xlarge	55	30	$0.84 / hour
c3.8xlarge	108	60	$1.68 / hour

Table 1: A Subset of AWS EC2 Instance Types and Pricing [5]

1. Introduction

Cloud computing shifts computing from local dedicated resources to distributed, virtual, elastic, multi-tenant resources. This paradigm provides end-users with on-demand access to computing, storage, and software services [1]. A number of cloud computing providers, such as Amazon Web Services (AWS) [2] and Google Compute Engine (GCE) [3], offer cloud computing platforms that provide custom applications with high availability and scalability. Users can allocate, execute, and terminate the instances (i.e., cloud servers) as needed, and pay for the cost of time and storage that active instances use based on a utility cost model [4].

To satisfy the computing resource needs of a wide variety of application types, cloud providers offer a menu of server types with different configurations of CPU capacity, memory, network capacity, disk I/O performance, and disk storage size. Table 1 shows a subset of the server configurations provided by AWS as of September 2014. For example, the *m3.medium* server with 3 CPU and 3.75 GB memory costs $\$0.07/hour$, while the more powerful *m3.2xlarge* server costs $\$0.56/hour$. By comparison, the GCE cloud provider offers competitive options, as shown in Table 2. Although the server types are named differently, the resource configurations at each pricing range vary only slightly. Cloud computing users must use this information to determine the appropriate subset of these resource configurations that will run an application cost-effectively, yet still meet its QoS goals, such as response time.

A common use case of the cloud is to offer existing software products, particularly web-based applications, through a *software as a service* (SaaS) model. In a SaaS model, the SaaS provider runs the web application in the cloud and customers remotely access the software platform, while the

Type	Virtual Cores	Memory (GB)	Price
n1-standard-1	1	3.75	\$0.07 / hour
n1-standard-2	2	7.5	\$0.14 / hour
n1-standard-4	4	15	\$0.28 / hour
n1-standard-8	8	30	\$0.56 / hour
n1-standard-16	16	60	\$1.12 / hour

Table 2: A Subset of GCE Machine Types and Pricing [3]

provider manages and maintains the software. SaaS providers typically provide service level agreements (SLAs) [6] to their clients, which dictate the number of users they will support, the availability of the service, the response time, and other parameters. For example, a provider of a SaaS electronic medical records system may offer an SLA that ensures a certain number of employees in a hospital can simultaneously access the system and that it will provide response times under 1 second.

An important consideration of SaaS providers is minimizing their operational costs while ensuring that the QoS requirements specified in their SLAs are met. For example, in the medical records system outlined above, the SaaS provider would like to minimize the cloud resources allocated to it to reduce operational cost, while ensuring that the chosen cloud resources can support the number of simultaneous clients and response times agreed to in the client SLAs. Moreover, as new clients are added and QoS requirements become more stringent (particularly in terms of the number of supported clients), the SaaS provider would like to know how adding resources on-demand (known as "auto-scaling") will affect the application's performance. Blindly allocating resources to the application to meet increasing load is not cost effective. Auto-scaling should instead be guided by a firm understanding of how resource allocations impact QoS goals.

Open Problem. While conventional cloud providers support simple and relatively quick resource allocation for applications, it is not an easy or straightforward task to decide an optimized and cost-effective resource configuration to run a specific application based on its QoS requirements. For instance, if a custom web application is expected to support $1,000$ simultaneous users with a throughput of $1,000 requests/minute$, it is hard to decide the type and minimum number of servers and cloud providers needed by simply looking at the hardware configurations.

Cloud providers and users have traditionally performed complex experimentation and load testing with applications on a wide variety of resource configurations. The most common practice is to deploy the application and perform a load stress test on each type of resource configuration, followed by analysis of the test results and selection of a resource configuration [7]. A number of load testing tools (e.g., jMeter [8, 9], ApacheBench [10], and HP LoadRunner [11]) are available to trigger a large amount of test requests automatically and collect the performance data. The situation gets even harder when an application includes a complex distributed architecture with multiple tiers, which requires the allocation of different groups of resources for the multiple tiers, followed by linking and connecting them properly.

Despite the importance of selecting a cost optimized resource allocation to meet QoS goals, many organizations do not have the time, resources, or experience to derive and perform a myriad of load testing experiments on a wide variety of resource types. Instead, developers typically employ a trial and error approach where they guess at the appropriate resource allocation, load test the application, and then accept it if the performance is at or above QoS goals. Optimization is usually only performed months or years into the application's life in the cloud, when insight into the affect of resource allocations on QoS goals is better understood. Even then, the optimization is often not systematic.

The primary challenges that impede early resource allocation optimization stem from limitations with load testing tools and a number of manual procedures required in the cloud resource optimization process. It is often hard to specify the customized load tests and correlate load test configuration with the expected QoS goals. Moreover, it is tedious and error-prone to manually perform load testing with different cloud resources to derive an optimized resource configuration, particularly when the application has a complex multi-tier architecture in a heterogeneious cloud platform.

Even when an optimized resource configuration is finally obtained, allocating and deploying all of the resources also requires significant orchestration and complexity. Prior research has addressed some challenges separately (e.g., modeling realistic user test behavior to produce customized load test [12], monitoring target test server performance metrics for capacity planning [13]). However, a comprehensive, fully automated approach designed specifically for benchmarking, deriving, and implementing optimized cloud resource allocations has not been developed, particularly for multi-tier applications.

Solution Approach ⇒ QoS-Oriented Modeling Framework for

Resource Optimization, Allocation and Recommendation System (ROAR). To address these challenges, this paper presents a model-based system called "ROAR" that raises the level of abstraction when performing load testing and automates cloud resource optimization and allocation to transparently convert users's application-specific QoS goals to a set of optimized resources running in the cloud. A textual domain-specific language (DSL) called the *Generic Resource Optimization for Web applications Language* (GROWL) is defined to specify the high-level and customizable load testing plan and QoS requirements without low-level configuration details, such as the number of threads to use, the concurrent clients, and the duration of keeping opened connections.

The model built from the GROWL DSL can generate a test specification that is compatible with our extended version of the jMeter load testing tool [8, 9]. Likewise, ROAR automates the process of deploying the application to the given cloud platform, executing the load test, collecting and aligning performance metrics, analyzing the test performance model, and controlling the next test iteration. ROAR can derive the appropriate cloud resource configurations to test, as well as automatically orchestrate the tests against each resource allocation before using the results to recommend a cost-optimized resource allocation to meet the QoS goals. When the optimal configuration is decided, ROAR can also generate the resource templates to automate the final deployment process.

The work presented in this paper extends the initial prototype of ROAR [14], with the following new contributions and extensions:

- Multi-tier distributed web applications that require deployment in different groups of cloud servers can be supported by ROAR. Multi-tier configuration can be specified in the GROWL DSL. Likewise, the ROAR controller can automatically allocate the right resources to deploy and link different tiers.

- The average latency of the deployment target has been added as another key QoS goal, in addition to the target throughput supported in our earlier work. The ROAR resource optimization engine aligns the average latency metrics together with the throughput and resource utilization metrics, and filters out resource configurations that support the target throughput but fail to meet an average latency goal.

- Rather than focusing on a single cloud platform (such as AWS), the

ROAR deployment manager has been extended to support deploying multi-tier applications to multiple cloud providers, which broadens its scope. In particular, we have included GCE as another platform supported by ROAR, which offers a wider range of optimization choices.

- The motivating example has been changed from a simple single-tier web application to a canonical three-tier web application with a database. Additional experiments and evaluation have also been included to showcase and quantify the performance, cost, and benchmarking effort for the resource configurations generated by ROAR.

The remainder of this paper is organized as follows: Section 2 summarizes a concrete motivating example in the context of a multi-tier web application, followed by describing the key challenges of cloud resource allocation and optimization in Section 3; Section 4 analyzes our solution by explaining each key component in the ROAR framework; Section 5 discusses the validation of this framework by presenting the generated resources for the motivating example; Section 6 compares our work on ROAR with the related research; and Section 7 presents concluding remarks.

2. Motivating Example

The example in this paper is based on a multi-tier web service called GhostBox built to support fine-grained control over mobile applications on the Android platform. GhostBox is designed to enable system administrators to configure specific security policies for employee smartphones and manage the policies remotely. These security policies specify the following properties:

- The types of the smartphone usage restrictions, such as running applications, accessing contacts, changing system settings and environments and
- The conditions to trigger the restrictions, which could be based on time, geographical locations, and context-aware signals (e.g., school WIFI, Bluetooth LE signals [15]).

By using GhostBox, administrators can use an administrative web portal to create and update security policies, while the smartphones used by employees periodically poll the server for updated security policies and implement the corresponding restrictions. For instance, administrators can create policies to disable gaming apps when a corporate WIFI signal is detected or enable sending/receiving text messages only from business contacts in the employers customer relationship management system.

The backend of GhostBox is built as a standard 3-tier web application. As shown in Figure 1, the presentation tier at the top-level interacts directly with the devices and admin users. It is implemented as a Tomcat web service, and exposes HTTP APIs that can be called by the GhostBox app running in the devices that poll for device usage policies.

GhostBox also integrates a web portal as a web service to provide a simple UI for admin users to configure the policies as well as manage all the users and group settings. The middle-tier is the logical tier which handles all the data access requests sent from the presentation tier. It runs a Jetty web service [16] and executes SQL statements to query and update the policies and other settings. A MySQL database runs in the bottom data tier to persist all the data related to policies, device usage statistics, etc.

From the perspective of the devices, the key functions provided by the GhostBox back-end are listed in Table 3. Likewise, Table 4 shows the QoS goals for each of the APIs. QoS is critical in this application since devices continuously poll the server and the service is sold based on a per-device

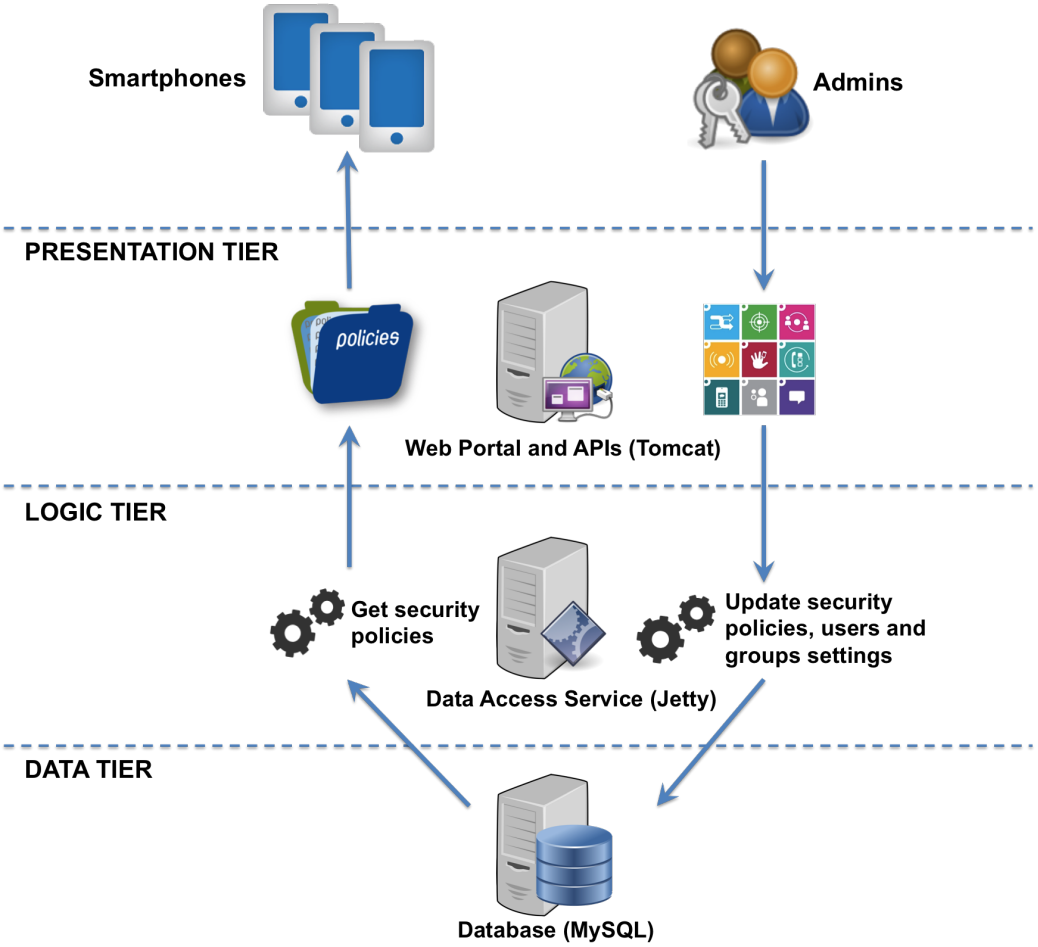

Figure 1: Overview of the GhostBox Multi-tier Architecture

Function	HTTP Method	URL
Get device policies	GET	/v1/policies/{deviceId}
Get policy config page	GET	/v1/policies/{deviceId}/view
Update policies	POST	/v1/policies/{deviceId}
Get settings config page	GET	/v1/account/{accountId}/view
Update settings	POST	/v1/account/{accountId}

Table 3: Key HTTP APIs Provided by GhostBox Backend

Function	Expected Frequency	Throughput	Latency
Get device policies	1 / 5 secs	24000 / min	100ms
Get policy config page	1 / 8 hours	4.1 / min	100ms
Update policies	1 / 8 hours	4.1 / min	100ms
Get settings config page	1 / 48 hours	0.7 / min	100ms
Update settings	1 / 48 hours	0.7 / min	100ms

Table 4: The QoS Requirements of the Key HTTP APIs

monthly cost. The SLAs for the service specify the maximum time from when a policy is enabled by an administrator until it reaches all devices with a network connection.

Getting the device policies is the most frequently called API by all the individual devices. Each device polls for new policies every 5 seconds so that it can make sure updated policies are retrieved and take affect immediately. Managing the policies and accounts is done in the web portal. Admin users can login to the portal and build policy configurations. Making a change involves getting the web page resources loaded first, followed by sending an HTTP POST request to send a new policy to the server. We assume that an admin updates the policies 3 times per day, while the account settings are updated in a less frequent rate - once every 2 days.

Assume that GhostBox needs to be provisioned to support 2,000 users. Based on the frequency of calls to each API from the client, it is possible to predict the minimum throughput needed to support the clients, as shown in Table 4. In addition to controlling throughput, another key QoS requirement related with the success of the application is to keep latency below a threshold. The maximum allowable latency for each function is given in the last column of Table 4.

A key requirement, therefore, is determining how to appropriately provision resources to meet the QoS requirements of the application. Determining how to provision cloud resources for GhostBox is non-trivial since it is based on a 3-tier distributed application architecture that is hard to model and characterize with existing analysis-based approaches.

3. Challenges

Determining the most cost-effective set of resources to meet a web applications QoS requirements without over-provisioning resources is a common problem in launching services in production or transitioning services to a SaaS model. The conventional practice is to benchmark web application performance using load testing on each set of possible target cloud resources. Developers use the results from these tests to estimate the required resources needed by the application to meet its QoS goals. This section summarizes the key challenges faced when deriving cost-optimized cloud resource configurations for web applications.

Challenge 1: Translating high-level QoS goals into low-level load generator configuration abstractions. To produce the huge amount of test traffic needed for QoS metrics, most tools use a multi-threaded architecture for load testing. The load test modeling abstractions from these tools are focused on the low-level threading details of the tooling rather than higher-level QoS performance specifications. For instance, jMeter, one of the most popular load testing tools, requires the specification of Thread Groups including the number of threads, the ramp-up period, and loop count.

Wrk [17] takes command-line parameters to control the number of connections to keep open, the number of threads, and the duration of the test. All these OS-level details on threads create an additional level of complexity for end-users, since they must translate the desired QoS goals to test, such as a throughput target, into the right threading configurations that can generate the appropriate load profiles to determine the feasibility of these targets. Moreover, increasing the number of threads does not always increase the throughput linearly [18], which makes the translation of high-level QoS goals into low-level load testing tool configuration abstractions even harder. Developers traditionally have manually analyzed their needs to derive the appropriate threading and ramp up configurations, which is non-trivial in many cases.

Challenge 2: Resource bottleneck analysis is challenging because current tools do not collect or correlate this information with QoS metrics from tests. It is essential to understand resource utilization of allocated cloud resources in order to identify bottlenecks and make appropriate resource allocation decisions. Moreover, temporal correlation of these resource utilizations with QoS performance metrics throughout the tests is essential under the following conditions:

- When required QoS goals are not met, identifying the resource bottlenecks in order to adjust resource allocations (e.g., add more memory or switch to a more powerful CPU) is difficult
- Even if the required QoS goals are satisfied, deriving resource utilization in order to estimate and ensure resource slack to handle load fluctuation (e.g., ensure 20% spare CPU cycles) is hard
- When QoS goals are satisfied, ensuring that there is not substantial excess capacity (e.g., CPU utilization is at 70% or less) is challenging
- For multi-tier web applications, the bottleneck may only be present in certain tier(s) and it is therefore important to understand the accurate resource utilization status in each individual tier and its relationship with the overall QoS or a request.

The goal is to find the exact resource configuration for each application tier where the QoS goals are met, there is sufficient resource slack for absorbing load fluctuations, there is not too much excess capacity, and there is no more efficient resource configuration that better fits the QoS and cost goals. Manual monitoring, collection, and temporal correlation of QoS metrics with resource utilization data is tedious and often inaccurate.

Challenge 3: Increased benchmarking complexity presented in large-scale complex web applications. The resource configuration design space for an application has a huge number of permutations. A key task in optimizing resource allocations is selecting the appropriate points in the design space to sample and test in order to make resource allocation decisions. Moreover, developers must decide when enough data has been collected to stop sampling the configuration space and make a resource allocation recommendation.

To find an optimized set of cloud resources for running the application with the desired QoS, the same load test must be run multiple times against the different samples of cloud resources. By comparing the different test data and performance metrics, the final decision on the required cloud resources can be made. For a single tier web application, which can be deployed on a single host, it is straightforward to benchmark and compare the performance using different resources (e.g., different server types).

When an application involves distributing its components and tiers in different environments or hosts, however, the benchmarking complexity and

effort can increase exponentially. For instance, to completely benchmark and compare the performance for the 3-tier web application in the motiving example, there will be a total of ${}_nC_3$ possible resource allocation configurations, where n is a the total number of available server types. If we want to analyze all of the 23 EC2 instance types provided by AWS, 1,771 load tests have to be performed.

Challenge 4: Lack of end-to-end test orchestration of resource allocation, load generation, resource utilization metric collection, and QoS metric tracking. Current load testing tools only focus on the generation of loads and the tracking of a few QoS metrics. However, other key aspects of a test, such as the automated allocation of different resource configurations sampled from the resource configuration space or the collection of resource utilization data from the allocated resources, are not managed by the tools. Allocating cloud resources for complex web applications such as multi-tier web architecture requires a number of system-level configurations (e.g., security groups, load balancers, DNS names, databases) that current tools force developers to manually manage.

Although most cloud providers offer tools to automate the allocation and deployment of cloud resources, these tools require manual manipulation of many low-level details to configure. For example, Cloud Formation, [19] provided by Amazon AWS, is an effective tool to automatically deploy an entire server stack with the cloud resources specified. However, it is a JSON-based specification that includes the configuration of over 50 resource types and hundreds of parameters, and is completely disconnected from current load testing tools. Moreover, once the resource allocations are made and tests launched, the resource utilizations on each cloud server need to be carefully remotely tracked, collected, and temporally correlated with QoS metrics tracked in the load testing tool.

Challenge 5: Extra configuration effort is required to evaluate and utilize different cloud providers. With the public cloud computing market getting increasingly popular, a number of key cloud providers offer very compelling resources and services to compete with each other. While the competition provides customers a broader range of choices on cloud resources that helps further optimize the resource configurations and pricing, extra effort is needed to carry out the benchmark and analysis across different cloud platforms, due to the different infrastructures, service features, management systems, software SDKs, etc.

For instance, accessing an AWS EC2 instance usually requires standard

SSH communication with security keys, while GCE uses its own tooling and commands to login to a remote server based on OAuth2 [20]. Although both AWS and GCE provide a deployment template mechanism for resource allocation in JSON (the AWS Cloud Formation and Google Cloud Deployment Manager [21]) the format and supported type of resources vary dramatically. Handling the diversity of the cloud platforms can therefore ben an obstacle to utilize all the compelling cloud resources and make the most optimized decision.

In summary, a gap in research exists between the current load testing techniques and the resource allocation problems of cloud resource optimization. In particular, load testing aims to verify if the given resources are sufficient to meet the requirements. Conversely, the resource allocation tries to figure out exactly what are the most optimized (i.e., most cost-effective) resources to meet the requirements. In the context of multi-tier web applications and different cloud providers, this gap becomes even more substantial.

4. Solution: Resource Optimization, Allocation and Recommendation System (ROAR)

To address the five challenges presented in Section 3 associated with load testing to direct and automate resource allocation and optimization, we developed the *Resource Optimization, Allocation and Recommendation System* (ROAR). ROAR combines modeling, model analysis, test automation, code generation, and optimization techniques to simplify and optimize the derivation of cloud resource sets that will meet a web application's QoS goals. This section explains each key component in the ROAR framework, which is summarized in Figure 2.

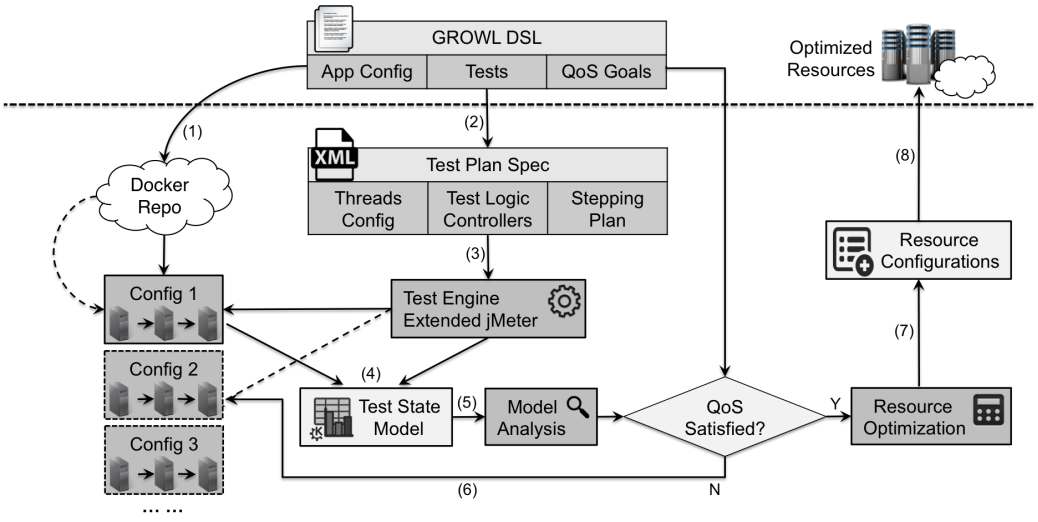

Figure 2: Overview of the Resource Optimization, Allocation and Recommendation System (ROAR)

At the heart of ROAR is a high-level textual DSL called the *Generic Resource Optimization for Web applications Language* (GROWL). ROAR uses GROWL to capture key resource configuration space and QoS goal information that is not currently captured in existing load generation tools (*Addressing Challenge 1*). The GROWL specification is translated into specific test plan inputs to the ROAR controller. ROAR then provides a fully automated end-to-end benchmarking and optimization orchestration, which handles resource allocation, application deployment, load test execution, performance metrics collection, performance model analysis, and test termination (*Addressing Challenge 4*).

ROAR can monitor detailed utilization metrics of cloud resources and align these metrics with the QoS performance metrics (*Addressing Challenge 2*). The aligned utilization/QoS metrics are then injected into a temporal performance test state model. This model can be used to perform model analysis and determine the next testing steps (*Addressing Challenge 3*).

ROAR supports conducting the same testing orchestration across different cloud platforms (e.g., AWS and GCE) based on the specification given in GROWL. When a final optimized resource configuration is determined, ROAR uses code generation to generate platform-specific deployment templates so that the allocation of the resource configurations for the chosen cloud provider can be automated (*Addressing Challenge 5*).

4.1. High-Level DSL Aligned with the QoS Domain

Based on our experience using a variety of load testing tools, we determined the following elements of their typical abstractions that were not properly aligned with the domain of cloud resource allocation derivation:

- The low-level test thread group configurations (e.g., the number of threads, thread ramp-up time, duration of open connections)
- The complex test logic controllers based on procedural programming concepts
- The inability to specify test application deployment and resource configurations, as well as the target QoS goals.

To address these abstraction issues, we developed the GROWL DSL to enable developers to specify web application resource configurations associated with QoS goals. GROWL will be used as the input to ROAR to automatically generate platform-specific load testing plans without requiring users to understand the underlying low-level jMeter-specific details, such as thread group or logic controllers.

```
apps {
    platform : aws;
    app {
        id : web-tier;
        containerId : roar-test/web-tier;
        port : 8080;
        isTestEndpoint : true;
        env : {
            "LOGICAL_ENDPOINT" : "{roar-test/logical-tier}"
        }
    }
    app {
```

```
        id : logical-tier;
        containerId : roar-test/logical-tier;
        port : 8900;
        env : {
            "DB_ENDPOINT" : "{roar-test/db-tier}"
        }
    }
    app {
        id : db-tier;
        containerId : roar-test/db-tier;
        port : 3306;
        isTestEndpoint : false;
    }
}

tests {
    healthCheckUrl : "/v1/ping";
    ordered : false;
    sampler {
        method : GET;
        path : "/v1/policies/testid";
        percentage : 99.9;
    }
    sampler {
        method : GET;
        path : "/v1/policies/testid/view";
        percentage : 0.01;
    }
    sampler {
        method : POST;
        path : "/v1/policies/testid";
        percentage : 0.01;
        requestBody : "{\"deviceId\":\"testid\", \"policies\":\"...\"]}";
    }
    sampler {
        method : GET;
        path : "/v1/account/testaccount/view";
        percentage : 0.01;
    }
    sampler {
        method : POST;
        path : "/v1/account/testaccount";
        percentage : 0.01;
        requestBody : "{\"accountId\":testaccount, \"settings\":\"..."]}";
    }
}

performance {
    throughput : 24050;
    timeunit : MINUTE;
    latency : 50;
}
```

Listing 1: Test Configuration Example in GROWL

Listing 1 shows the GROWL specification of the test plan for our motivating example, which contains the three major sections discussed below.

The apps section.. This section describes the basic information about each of the application components (e.g., tiers) contained in the application stack to test. The attribute *platform* specifies the cloud platform to use, which could be either *aws* (i.e., Amazon Web Services) or *gce* (i.e., Google Com-

puting Engine). A list of *app* structures describe the configurations for each application component via the following attributes:

- *id* is the identification of the application component, which is required and needs to be unique.
- *containerId* is a concept from Docker [22] that specifies the ID of the target deployable and runnable web application container stored in Docker repository [23]. To automate the application deployment to different target servers for testing, we used a containerization approach, based on the Docker container mechanism, to ease the process. Docker and the automated deployment process are discussed further in Section 4.2.
- *port* is the port number that will be used by the web application.
- *isTestEndpoint* indicates whether the application component will expose the public endpoints for the testing. Setting this attribute to *true* means that this application will be the top-tier to face users traffic directly;
- *env* specifies the environment variables for the operating system to be used in the allocated resources. ROAR uses this environment variable to connect different application components or tiers by letting others know how multiple applications should depend on each other. In Listing 1 an application reference (the string surrounded by the parenthesis) is used to present the actual value for the endpoint, which will be translated into the physical IP address or public DNS name in the load testing stage. The preferred multi-tier configuration and environment variable setup will be explained in Section 4.2.

The tests section.. This section configures the details on what HTTP APIs or web resources to test and how to test them. The *tests* section contains a list of test samplers. Each *sampler* represents a specific type of HTTP request that we want to include in the QoS specification. The typical HTTP request configurations, such as HTTP Method (e.g., GET, POST, DELETE), HTTP URI, HTTP Headers/Parameters, HTTP Request Body can be provided for each sampler. The *percentage* attribute in each sampler controls the frequency of each sampler in the overall test plan.

For example, the percentage 99.9 means that 99.9% of all the requests in the load test will go to this API path. This attribute is the key to simplify the usage of traditional jMeter logic controllers. By default, all samplers will be executed in random order with the given percentage. However, if the *ordered* attribute in tests is set to *true*, each sampler will be executed in the order as specified in the *tests* section. The other attribute *healthCheckUrl* is not used during the load test. Instead, the ROAR test engine relies on the URL specified in this field to check if the application to test is running correctly.

The performance section.. This section specifies the expected QoS goals we want to support in terms of both *throughput* and *latency*. The example in Listing 1 shows a throughput requirement for $24050 requests/minute$ with a maximum latency to be $100ms$. GROWL has been implemented using xText [24]. An XML-based jMeter compatible specification will be generated from a GROWL model instance as the input to the ROAR controller through xTend [25].

4.2. Testing Resource Allocation and Application Deployment

The only thing developers using ROAR need to do manually is create the GROWL model for the test plan - the rest of the testing and resource derivation process is automated for them. As shown in Figure 2, the basic workflow of the automated end-to-end test process is to first run the web application in the most resource constrained configuration, followed by the iterative benchmarking, analysis, and optimization processes. However, deploying the multi-tier distributed web applications across the various cloud resources is not an easy task, due to the following challenges:

- Different applications require setting up different software libraries, environment configurations and even operating systems, which makes it hard to build a generic automated deployment engine.
- When trying to support deploying the same application stack to different cloud providers, the deployment mechanism needs to be adaptive to the heterogeneous cloud platforms (e.g., SDKs, authentication mechanism, configuration parameters, etc.).
- When the distributed application components depend on each other, the extra complexity is added to handle their links and correct communication.

To solve these problems, we built a container-based deployment manager as part of ROAR that can support generic application stack deployment in various cloud platforms. The process-based container tool Docker [26], can package an application and its dependencies in a virtual process-based container that can run on any Linux server. Users only need to package their application in a Docker container with the preconfigured execution environment and push the application to the centralized Docker container repository.

The ROAR deployment manager will pull the application container based on the *containerId* specified in GROWL and run it automatically (Step (1) in Figure 2). With the flexibility and portability of Docker containers, we can deploy a web application in any type of the target cloud server at anytime. The server does not need any special environment configuration except installing Docker itself, which has in fact already been included in most of the cloud VM images by the major cloud providers [27][28].

Another important capability provided by Docker is the ability to version containers in Git. Each change to a container can therefore be tracked, which enhances the repeatability and traceability of deployment and configuration changes. GROWL models and the generated test performance state models can also be versioned along with Docker containers to understand how the performance and resource allocations change as the application evolves.

To connect the distributed application components for all the tiers, knowing the endpoint of each component is indispensable. In ROAR, we follow the common practice in distributed systems to configure the endpoint through environment variables, with the main benefit being that it is easy to change the value without affecting the built software artifact, and it is simple to be integrated with the deployment mechanisms (e.g., deployment scripts). As shown in Listing 1, a set of environment variables can be configured as key/value pairs in GROWL. The actual physical endpoint is not known until the cloud server has been allocated and is running. The application ID can therefore be used as a reference so that the deployment manager knows their dependency relationship, and the reference will be replaced with the final address when the cloud server is running.

The cloud resource allocation and management is done through the SDKs or APIs provided by the cloud providers. Currently, ROAR supports both AWS and GCE. When you specify the *platform* in GROWL, the ROAR deployment manager switches to use the correspondent implementation to communicate with the target platform to request the cloud resources. It first waits for the specific cloud server type to startup. When it is ready,

the Docker command is triggered to pull the application container, followed by initializing the environment variables in the container with the actual IP addresses or DNS names associated with the instances. Different tiers will automatically discover each other based on the endpoint of the environment variables. Figure 3 shows the capability of the ROAR deployment manager in this context.

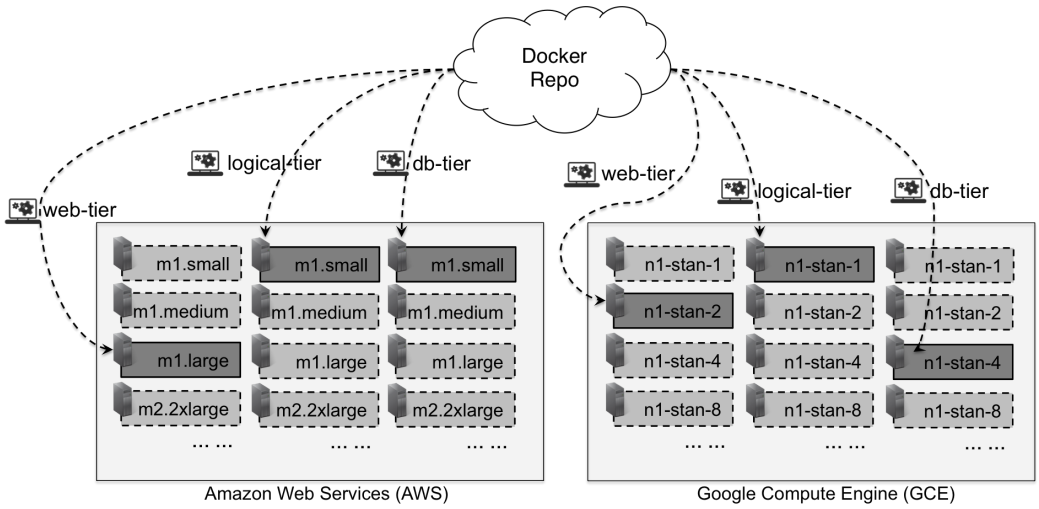

Figure 3: ROAR's Deployment Manager Pulls Applications from Docker and Deploys Them to Servers in Different Cloud Platforms

4.3. Load Test Execution

After the entire application stack has been deployed, the ROAR test engine periodically checks the running status of the application with the *healthCheckUrl* before triggering the load test. In ROAR, jMeter is chosen to be the core test engine due to its support on a wide range of web applications, as well as its extensibility.

jMeter supports an XML-based specification for the load test plan. A GROWL interpreter has been developed to generate the jMeter specification from the GROWL model (Step (2) in Figure 2). One part of the generated jMeter test specification, based on the *tests* section in GROWL, is the test plan, which uses a series of jMeter logic controllers to construct a test that loads the various web application paths according to the percentages or ordering specified in GROWL.

For example, the two UPDATE APIs are not called very often, so the generated jMeter test plan uses an *OnceOnlyController* to send only one request for the entire test period. The GET APIs are put inside of a *ThroughputController* provided by jMeter to accurately control the execution frequency of their samplers. Finally, all four controllers are nested inside an *InterleaveController* that indicates to jMeter to alternate between each of the controllers in each loop iteration. Clearly, there is a significant amount of logic and associated abstractions that are specific to jMeter and disconnected from the domain of specifying desired web application QoS.

The standard jMeter test plan only allows developers to specify the number of concurrent threads to use and provides no guarantee that the test will sufficiently load the web application to assess if the QoS goal is met. If the number of threads configured is not sufficient, jMeter will not be able to test the full capacity of the target server. However, if the number configured is too large and they ramp up too quickly, the server may run out of CPU or memory before exposing the actual throughput and server resource utilization under load. Developers must therefore rely on trial and error specification of these parameters to test each target throughput.

To overcome this issue, we developed a model analysis to determine the appropriate test strategy to assess the target throughput goal. To produce test loads with the desired throughput accurately, we built a customized jMeter plugin [13] and throughput shaper that analyzes and operates on the GROWL model. The shaper is capable of dynamically controlling the number of threads being used to ensure that the test meets the GROWL goals. Based on the needed throughput, it automatically increases or decreases the number of threads to reach the throughput accurately, while simultaneously watching the resource utilization of the target server to ensure that it isn't overloaded too quickly.

Our throughput shaper supports gradual throughput increments. A throughput stepping plan is automatically derived from the specified throughput QoS goals in GROWL. For instance, if the target throughput goal is $5,000 requests/minute$, the generated test plan divides the test into 10 different stepping periods, so that the test load increases from 0 to $500 requests/minute$ within the first 5 seconds and then keeps running at $500 requests/minute$ for another 20 seconds before jumping to the next level at $1,000 requests/minute$.

The test load is gradually increased, rather than maximizing the load from the start, in order to avoid crashing the server if the load is substantially more than the capacity of the server. In addition, the incremental ramping periods

provide more diverse sampling metrics used in the test model analysis in the next step. The test plan specification is used by the extended jMeter engine to trigger the load test (Step (3) shown in Figure 2).

4.4. Test Performance Data Collection and Alignment

The key function of jMeter is to collect, aggregate, and summarize the QoS performance metrics for the applications being tested, such as throughput, latency, and error rate. What is missing from jMeter, however, is the capability of monitoring and collecting resource utilization metrics from the servers used for hosting the application. A key component of ROAR is thus a mechanism to collect and correlate server resource utilization with QoS performance metrics.

The ROAR throughput shaper discretizes the test epoch into a series of discrete time units with known QoS values. The QoS values at each discrete time unit are correlated with server resource utilization during the same time unit. To record the actual throughput in each moment, we modified the native jMeter *SimpleReporter* listener to be capable of automatically reporting the throughput every single second (the native *SimpleReporter* only reports the current aggregated throughput). Likewise, to record the target server performance metrics, another customized jMeter plugin [13] is applied, which runs a web-based performance monitor agent in the target server. HTTP calls can be made from jMeter to retrieve the performance metrics at runtime and assign them to the appropriate time unit for correlation with QoS values.

A range of detailed resource utilization metrics are supported by our plugins. However, we only record CPU, memory, network I/O and disk I/O, since we have found these are the primary resource utilization metrics for making resource allocation decisions. The extensions we made to the jMeter plugin include (1) using a global clock that is consistent with the time used by the QoS performance records, (2) enabling the collection process in non-GUI jMeter mode, and (3) distributed resource utilization collection for each application using the global clock.

Based on the global clock, the collected resource utilization and QoS metrics are aligned and stored in a test state model (Step (4) in Figure 2). The test state model is the input to the bottleneck analysis and resource optimization algorithms run later and can also be loaded in the performance viewer to display the results visually. As shown in Figure 4, based on the

global clock, the correlation of QoS (throughput and latency) and resource utilization (CPU/Memory/Network/Disk utilization) can be clearly seen.[1]

4.5. Model Analysis and Termination Strategy

The test state model containing the global clock-aligned QoS values and resource utilizations allows us to analyze resource bottlenecks and derive the best cloud resource configurations to meet the modeled QoS goals. One issue that must be addressed before resource configuration derivation is potential points of instability in the data, which show an inconsistent correlation between the QoS and the resource utilization with the values recorded during the rest of the test epoch. These periods of instability are due to jMeter's creation and termination of a larger number of threads in the beginning and end of each testing step triggered by the throughput shaper.

4.6. Handling Performance Instability

To handle the instability, a short 10-second warmup period has been automatically added in the beginning of each test plan. The $1/10^t h$ of the target throughput load will be used during the warmup and the collected metrics during this period of time will be ignored. After the warmup, the test plan gradually increase the test loads by $1/10^t h$ of the target throughput and keeps running it for 20 seconds. Based on our experiments, the performance and utilization metrics usually fluctuate during the first half of each individual step, which does not reflect the real QoS that can be sustained by the application at that load. Thus, the final test performance model only use the second half of the data to average the values for both QoS performance and resource utilization metrics.

Therefore, if n is the number of total steps used in the test plan, after load testing, a series of aligned data set on the actual throughput (T), average latency (L), CPU utilization rate (U_{cpu}), memory utilization rate (U_{mem}), network utilization rate (U_{net}) and disk utilization rate (U_{disk}) will be available in the following test performance model:

[1] Disk I/O is not involved with this application.

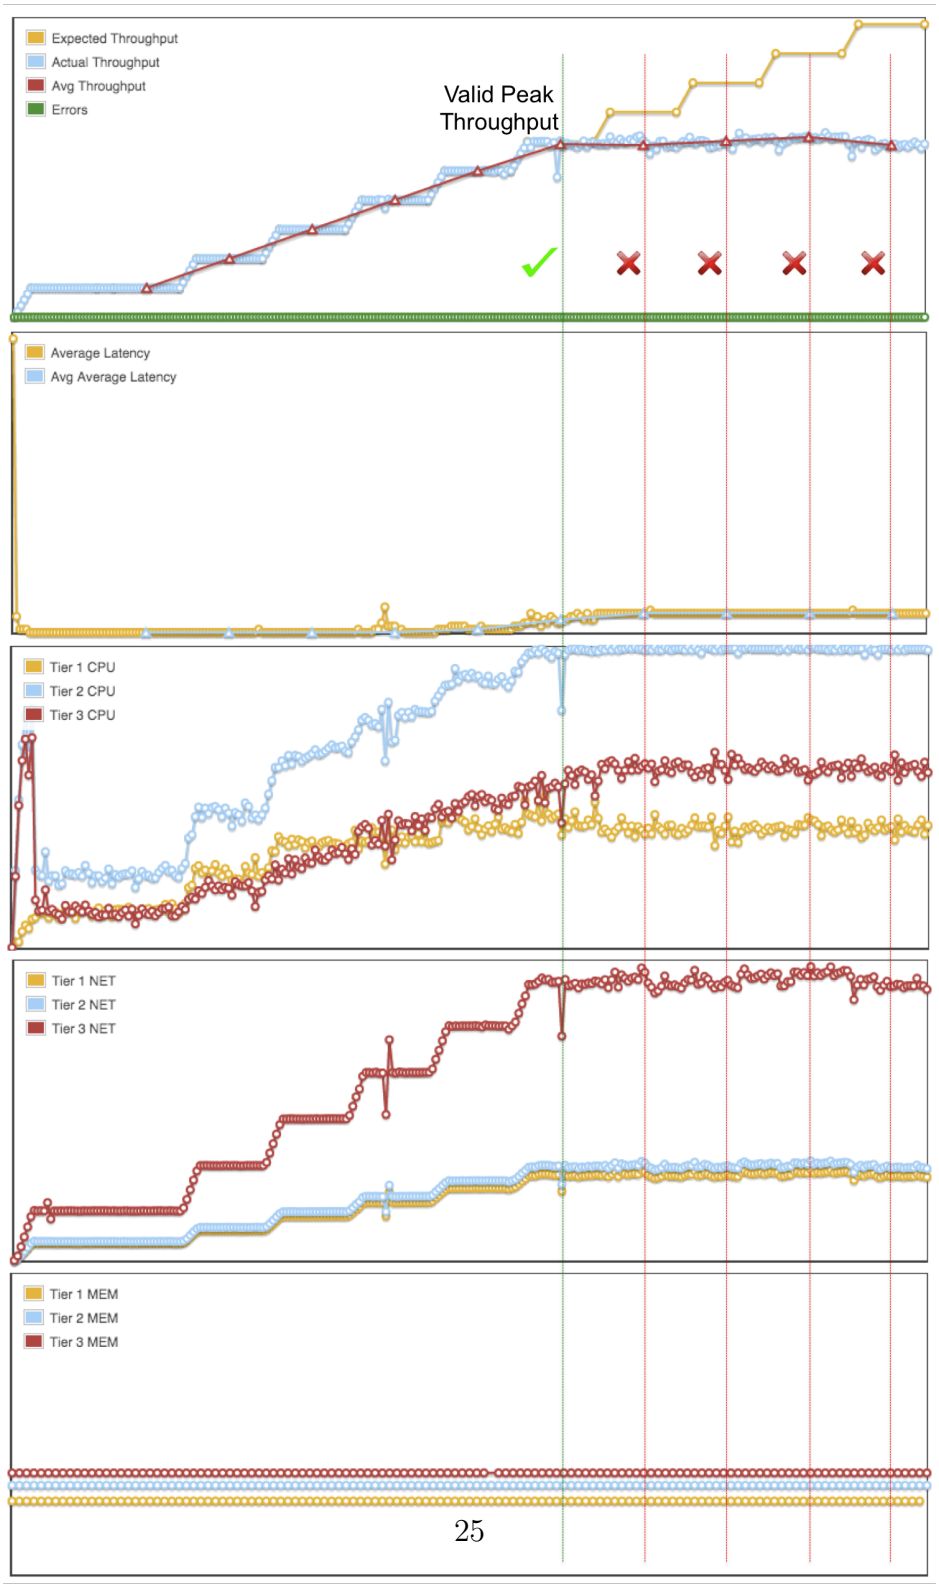

Figure 4: A Sample Test State Model with Aligned Data on Throughput, Average Latency, CPU, Memory, and Network

$$
\begin{aligned}
T &= \{t_i | 1 \leq i \leq n\} \\
L &= \{l_i | 1 \leq i \leq n\} \\
U_{cpu} &= \{u_{cpu_i} | 1 \leq i \leq n\} \\
U_{mem} &= \{u_{mem_i} | 1 \leq i \leq n\} \\
U_{net} &= \{u_{net_i} | 1 \leq i \leq n\} \\
U_{disk} &= \{u_{disk_i} | 1 \leq i \leq n\}
\end{aligned}
$$

4.7. Resource Allocation Analysis

The resource allocation analysis can be performed with a variety of models and optimization goals [29, 30, 31]. ROAR uses an optimization engine to select the minimum number of servers N needed to run each application tier to reach the expected throughput within the expected latency.

Let T_{ex} be the expected throughput and L_{ex} be the expected maximum latency specified in GROWL. If T_p is the actual peak throughput with a certain resource configuration, then:

$$N = \lceil T_{ex}/T_p \rceil \tag{1}$$

The key problem here is to find a reasonable T_p from all the data points in the test state model without violating the expected latency goal T_{ex} and the resource utilization limits. Generally, when the peak throughput is reached, one or more resource utilizations will also approach 100% utilization (e.g., CPU or memory utilization goes to 100%, or network usage goes to 100%), and at the same time, the average latency of responses increases dramatically. We use these two conditions to determine T_p. Formally,

$$
\begin{aligned}
T_p &= \max_{0 \leq i \leq n} t_i \\
where, \quad & l_i \leq L_{ex} \\
& u_{cpu_i} < 100\% \\
& u_{mem_i} < 100\% \\
& u_{net_i} < 100\% \\
& u_{disk_i} < 100\%
\end{aligned}
$$

Figure 4 shows a typical dataset collected during a test. The throughput shaper produces an increasing number of requests as indicated by the expected throughput (yellow line in the top graph of Figure 4), but the actual throughput fails to meet the QoS goal after the 5th step (blue line in the top graph of Figure 4). Checking the performance metrics of the target server for each application tier, it can be found that although memory and network have not been fully utilized, the CPU utilization in Tier 2 (Data Tier) has reached its limit - 100%, and therefore the peak throughput using this server appears at the 6th step.

A complete test state model provides the detailed QoS performance with the given resource configuration for each tier. In order to compare and figure out the most optimized resource configuration, the same load testing process needs to be repeated for each type of resource to get the different values for N. The optimization engine in ROAR is capable of automating all the tests sequentially from the most resource constrained server types to more powerful and more expensive types (Step (6) in Figure 2).

In Amazon EC2, resources can only be allocated in fixed configurations of memory, CPU, and network capacity, rather than in arbitrary configurations. There are over 20 different instance types, which are virtual machines with fixed resource configurations, provided by AWS, so it is hard to manually sample the performance of all instance types. In addition, when the web application stack involves multi-tier applications like the motivating example, benchmarking all the possible combination of the server types with each tier would be tedious and costly.

After the test state model is generated, the ROAR optimization controller uses the test output data to choose the next resource configuration and then switches a tier to a bigger type of server with a more powerful hardware configuration if resource bottlenecks were found in that tier. To reduce the benchmarking effort, if the target QoS goals are met, the test can terminate after the next resource configuration is tested. The reason that the test can terminate is that any larger resource configurations will only have additional slack resources that are not needed to meet the desired QoS goals.

Conversely, when the target QoS goals are not met with a certain configuration, the ROAR test engine analyzes which application tier caused the bottleneck, i.e., the application tier with at least one type of resource utilization that reached 100%. ROAR will only change the instance type of that tier for the next test iteration until the QoS goals are met or other tiers become the bottlenecks. Using this approach, ROAR can effectively permute

the resource configuration of the application to efficiently find bottlenecks without exploring all possible resource configurations.

The entire testing process is repeated to derive a model of the QoS and resource utilization on the new resource configuration. This process is repeated until the most cost-effective configuration solution is found. A prototype of a web-based service has been implemented to take the GROWL as the input and generate the final cloud resource configurations ready for deployment.

The final resource configuration is decided by comparing and analyzing the results (Step (7) in Figure 2). The most cost-effective solution is derived by finding the minimum total cost of using each type of server to reach the desired throughput. Let I be the set of applications, and J be the set of available server types. If P_{ij} represents the cost of the server type j ($j \in J$) which is chosen to run application i ($i \in I$), the total cost is $C = Sum_{i \in I} P_{ij} * N_i$, and the most cost-effective solution is Min_C.

Besides finding the most cost-effective solution, knowing the solution that supports the fastest response time (i.e., the lowest response latency) is also essential for general web applications in practice. The fastest response time configuration is decided by the average latency measured for each type of server. Although the latency varies based on the testing environment (e.g., running the load test inside AWS would have a much lower latency than running the test outside of AWS), the comparison of the average latencies for different server types could produce a good indicator of the performance with which two different servers can handle the same request load. Thus, we choose the server type with the lowest average latency as the fastest response time configuration.

4.8. Automated Resource Template Generation

After the cloud resource allocation needed to meet the QoS goals is derived, ROAR can then automatically allocate the resources and deploy the web application. Manually allocating resources with a fleet of servers in a cloud environment is not an easy task in general. It's even harder when optimized service architecture patterns and best practices are needed.

For example, in AWS, the best practice of deploying multiple servers to support the same web application is to use an Elastic Load Balancer (ELB) [32] as the entry point. An ELB is capable of handling any number of requests from users and balancing the dispatch of the requests to each server behind it. A number of configuration parameters for security groups, private keys,

availability zones, server launch configuration data, and DNS names have to be specified correctly when using an ELB.

To ease the resource allocation process, ROAR uses a generative approach to produce the configuration artifacts needed for each cloud provider. Currently, ROAR supports artifact generation for both AWS and GCE, and support for OpenStack is planned. The AWS artifact generator is built on top of the AWS Cloud Formation service, which provides an engine for automating the deployment of multi-server architectures via templates defined in JSON. Based on the N and server types generated from the resource optimization engine, we fill the template with the appropriate configuration steps, resource types, and other parameters needed to launch the application.

After the template is generated, it can be uploaded to the AWS Cloud Formation console, allowing developers to deploy the web application in the derived configuration with a single click (Step (8) in Figure 2). A similar mechanism called the *Cloud Deployment Manager* (CDM) [21] is provided by GCE, which also uses JSON template but contains different configurations since the cloud resource types and management SDK varies dramatically. GCE also provides a load balancing service to include multiple servers and provides a single endpoint for each tier. The template generation is capable of filling the same type of cloud resource configuration in the platform-specific templates.

5. Evaluation

This section presents a case study evaluation based on the motivating example from Section 2. Based on the given web application stack, we run experiments with two sets of QoS goals in both cloud platforms, as shown in Table 5. These experiments allow us to observe the different optimizations provided by ROAR, as well as the difference between the various cloud computing platforms.

Test No.	Target Throughput (req/min)	Target Latency (ms)	Cloud Platform
1	24050	50	AWS
2	24050	50	GCE
3	50000	30	AWS
4	50000	30	GCE

Table 5: The List of Test QoS Goals and Platforms Used to Evaluate ROAR

By default, ROAR considers all the available instance types from each cloud platform. This experiment focuses on the general-purpose cloud servers that fits the nature of the motivating example. The goal is to clearly compare the differences in the same instance type category, as well as avoid the expense of the high-end server types (e.g., GPU or Cluster instance types, high memory or high I/O instance types). The instance types used from AWS include: *m3.medium, m3.large, m3.xlarge, m3.2xlarge, c3.4xlarge, c3.8xlarge*, while in GCE we choose: *n1-standard-1, n1-standard-2, n1-standard-4, n1-standard-8, n1-standard-16*, as listed in both Table 1 and Table 2. The chosen sets of servers in both platforms are very competitive in terms of both pricing and hardware configurations.

5.1. Resource Optimization Results

The test plan was specified using GROWL for each of the 4 tests to run, as shown in Listing 1. ROAR then automates the benchmarking and optimization process, with the results shown in Table 6 - 9. Each table lists the peak throughput and average latency data collected for each tier/instance type in different test iterations. Knowing the peak throughput, we can estimate the total number of servers required for each tier using the Equation 1 and then calculate the total cost. Using the minimum cost option for each tier (i.e.,

Tier	Instance Types	Peak Throughput (req/min)	Average Latency (ms)	Total Server Required	Total Cost (dollar/hr)
db	m3.medium	24050	11	1	0.07
db	m3.large	24050	10	1	0.14
data	m3.medium	6500	13	4	0.28
data	m3.large	19000	12	2	0.28
data	m3.xlarge	24050	10	1	0.28
data	m3.2xlarge	24050	10	1	0.56
web	m3.medium	8100	14	3	0.21
web	m3.large	20000	13	2	0.28
web	m3.xlarge	24050	13	1	0.56
web	m3.2xlarge	24050	11	1	0.56

Table 6: Resource benchmarking result for Test 1

Tier	Instance Types	Peak Throughput (req/min)	Average Latency (ms)	Total Server Required	Total Cost (dollar/hr)
db	n1-standard-1	24050	12	1	0.07
db	n1-standard-2	24050	10	1	0.14
data	n1-standard-1	12000	17	3	0.21
data	n1-standard-2	17000	16	2	0.28
data	n1-standard-4	24050	14	1	0.28
data	n1-standard-8	24050	13	1	0.56
web	n1-standard-1	13000	17	2	0.14
web	n1-standard-2	18500	16	2	0.28
web	n1-standard-4	24050	16	1	0.56
web	n1-standard-8	24050	14	1	0.56

Table 7: Resource benchmarking result for Test 2

Tier	Instance Types	Peak Throughput (req/min)	Average Latency (ms)	Total Server Required	Total Cost (dollar/hr)
db	m3.medium	24500	33	3	0.21
db	m3.large	48000	28	2	0.28
db	m3.xlarge	50000	22	1	0.28
data	m3.medium	6500	38	8	0.56
data	m3.large	19000	37	3	0.42
data	m3.xlarge	28000	25	2	0.48
data	m3.2xlarge	46500	20	2	1.12
data	c3.4xlarge	50000	19	1	0.84
data	c3.8xlarge	50000	20	1	1.68
web	m3.medium	8100	39	7	0.49
web	m3.large	20000	30	3	0.42
web	m3.xlarge	31000	30	2	0.48
web	m3.2xlarge	50000	29	1	0.56
web	c3.4xlarge	50000	25	1	0.84

Table 8: Resource benchmarking result for Test 3

the gray row in each table), the most cost-effective resource allocation can be derived.

For example, when running *Test 1* in AWS, the best resource combination is 1 *m3.medium* instance for the *db-tier*, 1 *m3.xlarge* instance for *data-tier* and 3 *m3.medium* instances for the *web-tier*. Conversely, when running *Test 1* in GCE, the most optimized solution would be 1 *n1-standard-1* instance for the *db-tier*, 3 *n1-standard-1* instances for *data-tier* and 2 *n1-standard-1* instances for the *db-tier*. All of these critical pieces of information are derived automatically using GROWL and ROAR and would be difficult to obtain through current load testing practices with existing tools and approaches.

Most of the cloud providers precisely scale their server performance according to pricing. It is, therefore, likely that a configuration with multiple servers of smaller instance types can meet the same QoS goals as a configuration with fewer servers of bigger types. For instance, Table 6 shows that multiple *m3.medium* or multiple *m3.large* servers can reach the same performance as a single *m3.xlarge* for *data-tier* with the same cost. In this case, we often consider choosing the configuration with lower request latency.

Tier	Instance Types	Peak Throughput (req/min)	Average Latency (ms)	Total Server Required	Total Cost (dollar/hr)
db	n1-standard-1	27000	31	2	0.14
db	n1-standard-2	43000	28	2	0.28
db	n1-standard-4	50000	24	1	0.28
data	n1-standard-1	12000	40	8	0.56
data	n1-standard-2	17000	38	3	0.42
data	n1-standard-4	27000	25	2	0.48
data	n1-standard-8	46000	24	2	1.12
data	n1-standard-16	50000	24	1	0.84
web	n1-standard-1	12500	40	7	0.49
web	n1-standard-2	18000	38	3	0.42
web	n1-standard-4	29500	31	2	0.48
web	n1-standard-8	50000	23	1	0.56
web	n1-standard-16	50000	22	1	0.84

Table 9: Resource benchmarking result for Test 4

Conversely, if a fewer servers are preferred (e.g., to reduce operational and management complexity), the ROAR optimization engine can be configured to prioritize solutions with fewer server instances. For example, a single server for the database could be used to avoid a complex database clustering configuration. The results for *Test 3* and *Test 4* show that when the requirements on latency are changed to $30ms$, some server types are automatically filtered out when the benchmarked latency is above this threshold.

Regarding the optimized solution, we currently rely on the direct aggregation of the throughput results to estimate the maximum throughput when multiple instances of the same type are being used. In fact, a further optimization could be done in this space by mixing different types of servers. For instance, Table 8 shows a more cost-effective solution for the *web-tier* uses 2 *m3.large* plus 1 *m3.medium*. However, mixing different types of servers can create challenges in practice, because the load balancers equally distribute traffic by default, which means that smaller server types may receive the same amount of traffic as bigger ones and therefore the overall performance may not meet expectations.

5.2. Performance Validation

Satisfying all the QoS performance requirements is a key goal for any resource allocation task. It is essential to ensure that the generated cloud resource configuration by ROAR can actually handle the desired throughput and latency requirements. Although detailed benchmarking processes have been conducted for different configurations, the performance of aggregating all the applications and resources together is still based on estimation. To validate the final performance and QoS, therefore, performance experiments were done for all of the 4 test scenarios. Each test scenario was deployed using the generated resource configuration templates - load balancers with cloud servers behind them.

A dedicated load test was triggered for each test scenario based on the given QoS goals. By comparing the actual throughput and average latency to the ROAR estimated throughput and latency, we derived the average error rate of the throughput (e_{tp}) and latency (e_l) for each application in the experiment, as shown in Table 10. It can be seen that the estimates on throughput are within 0.4% of their actual values.

Test No.	Throughput (req/min)	e_{tp} (%)	Latency (ms)	e_l (%)
1	23950	0.4	24	0.0
2	24580	0.0	26	0.0
3	50100	0.0	29	0.0
4	49900	0.02	27	0.0

Table 10: Average Performance Deviation (%)

The 0.4% error for *Test 1* arises largely from the error of the throughput shaper used in the load test engine. For instance, when the provisioned throughput is $24050 requests/minute$, the load generated may not exactly reach $24050 requests/minute$ but instead top out at $23950 requests/minute$ due to thread contention, which produces a 0.4% error rate. Regarding the latency error rates, they are all 0%, meaning that the average latency constraints were always properly derived by the ROAR generated resource configurations.

The low error rates were expected because the throughput and resource allocations were based on actual load testing and benchmarking, with the only estimation being aggregating the peak throughput to support the target throughput.

Test No.	Number of Iterations to Reach QoS Goals	Maximum Number of Potential Iterations
1	10	444
2	10	300
3	14	444
4	13	300

Table 11: Benchmarking Effort Comparison

5.3. Benchmarking Effort

As described in Section 4.5, ROAR automatically scales up the type of instance used by an application tier when that tier reaches its utilization limit for one or more resources and terminates the test when the QoS goals are met. As given in Section 3, the total number of configuration combinations needed to test the motivating example would be 444 for *Test 1* and *Test 3* using 6 AWS instances, 300 for *Test 2* and *Test 4* using 5 GCE instances. By comparison, Table 11 shows the number of total benchmarking iterations required for each test. These results indicate that ROAR effectively reduced the benchmarking effort by using bottleneck analysis to navigate the resource allocation configuration space of each tier.

5.4. Cost of Resource Configurations

Knowing the most optimized resource allocation in each cloud provider for different QoS goals can provide users with critical guidance on cloud provider selection. Table 12 shows that the cost of using AWS and GCE are only slightly different based on the QoS goals. This slight difference is

Test No.	Cloud Platform	Total Cost (Dollars/Hour)
1	AWS	0.56
2	GCE	0.42
3	AWS	1.18
4	GCE	1.32

Table 12: Cost Comparison

expected because both cloud providers offer the same categories of servers at the same price ranges. Although their marketed hardware specification and parameters are different, their basic computing capabilities are very similar.

However, we did notice a performance difference existed between the two cloud providers. Particularly, for the motivating example web application used in this paper, the *n1-standard-1* instance from GCE always performed much better than the server in the same pricing range *m3.medium* from AWS. However, for the higher end servers in AWS, such as *m3.xlarge* and *m3.2xlarge*, they always showed better performance in terms of both throughput and latency than the equivalent GCE instance types.

6. Related Work

This section compares our work on ROAR with the following related research efforts.

Ferry et al. [33] summarized the state-of-the-art on cloud optimization and pointed out the need for model-driven engineering techniques and methods to aid provisioning, deployment, monitoring, and adaptation of multi-cloud systems. Revel8or [34] is a model-driven capacity planning toolsuite to solve the problems related to complex multi-tier applications with strict performance requirements. They use UML 2.0 to model and annotate design diagrams, and derive performance analysis models. However, Revel8or only supports applications developed in a model-driven approach from platform-independent model to platform-specific model to the final code, while our approach targets all web applications, including hand-written and legacy applications. TOSCA [35] is another DSL proposed by Binz et al. to specify the topology and orchestration configuration for cloud applications to ease the cloud management tasks. Unlike our work, however, it does not focus on resource optimization and allocation.

A number of research projects have focused on model-based load testing. Draheim et al. [12] used a modeling approach to produce realistic load testing scripts for web applications. Their approach simplifies the creation of realistic usage models of individual user behavior based on stochastic models. However, they do not analyze the minimum required cloud resources. Wang et al. [36] present a *Load Testing Automation Framework* (LTAF), which offers usage models to simulate users's behavior in load testing and workflow models to generate the realistic testing load. Similar to prior work, LTAF focused on simplified creation of realistic load testing configurations for web applications, with detailed performance reports. Since resource optimization is not their goal, neither of these two frameworks takes the target server resource consumption into consideration, in contrast to our work on ROAR.

Wei etl al. [30] applied game theory to solve cloud-based resource allocation problems using abstract resource and application mathematics models. Based on a specification of the application's QoS constraints, Binary Integer Mining and evolutionary algorithms are used to drive the optimization. This research differs from our work on ROAR in the following ways:

- it focuses on resource allocation for multiple applications being scheduled in large cloud environments and

- it emphasizes abstract analysis rather than empirical testing of production applications.

Li et al. [37] present a method for resource optimization in clouds by using performance models in the deployment and operations of the applications running in the cloud. An optimization algorithm is implemented to accommodate different goals, scopes and timescales of optimization actions, to minimize cost while meeting SLAs across a large variety of workloads. Recent work by Chalsiri et al. [38] focused on the cloud resource optimization problem from a different perspective. They examined provisioning algorithms to better utilize market-based resources that fluctuate in cost based on demand to reduce total cost. To handle the resource optimization for multiple applications, Frey et al. presented a search-based genetic algorithm to find the near-optimal solution that optimizes response times, costs, and number of SLA violations [39]. Their approach is based on the simulation tool called CDOSim to improve the search process, while our approach focuses on performing actual benchmarking to obtain the accurate data. Instead of using genetic algorithms, Catan et al. [40] specified application dependencies and criteria, as well as applied external constraint solvers to reach an optimal configuration. Although the framework they built can also handle multiple-tier web application, they focus on finding out the most optimized actions to transform the deployment configuration from one state to another when a reconfiguration request comes in, rather than targeting the QoS goals.

7. Concluding Remarks

This paper presented a QoS-Oriented modeling framework called the *Resource Optimization, Allocation and Recommendation System* (ROAR). ROAR automates the testing and derivation of optimized cloud resource allocations for web applications. In particular, ROAR automates the end-to-end test orchestration from application deployment, metrics collection and alignment, to test model generation and the test iteration with termination.

ROAR uses a textual DSL called GROWL to hide the low-level configuration and analysis details of load testing for the complex multi-tier web applications. QoS goals (including target throughput and latency) can be expressed via the DSL, which will be used to direct the generated load test plans. GROWL generates an optimized resource configuration to meet the web application's QoS goals, as well as a deployment and configuration template for the specific cloud provider.

Our future work focuses on enriching GROWL and exposing more best practices in load testing and resource optimization as simple DSL language constructs. Examples include supporting complex test behaviors that are driven by input parameters (e.g., generate random or variable parameters as test inputs) or time-based tests (e.g., send request r1 during the period of t1, and request r2 only during the period of t2).

Moreover, only average latency is currently supported as one of the QoS goals. Therefore, enabling other goals (such as percentile latency) would also be useful. While multiple-tier application can be automatically benchmarked and optimized, we assume that each tier is deployed separately in different servers. In some cases, co-locating the different application tiers together in the same host might offers a deeper optimization space. However, more sophisticated optimization algorithms would be needed to handle this case.

[1] B. T OGRAPH, Y. R. MORGENS, Cloud computing, Communications of the ACM 51 (7).

[2] Amazon Web Services (ELB), https://aws.amazon.com/ (2014).

[3] Google Compute Engine (GCE), https://cloud.google.com/products/compute-engine/ (2014).

[4] M. A. Rappa, The utility business model and the future of computing services, IBM Systems Journal 43 (1) (2004) 32–42.

[5] Amazon EC2 Pricing, http://aws.amazon.com/ec2/pricing/ (2014).

[6] Z. Shu, S. Meina, An architecture design of life cycle based sla management, in: Advanced Communication Technology (ICACT), 2010 The 12th International Conference on, Vol. 2, IEEE, 2010, pp. 1351–1355.

[7] D. Menascé, Load testing of web sites, Internet Computing, IEEE 6 (4) (2002) 70–74.

[8] E. H. Halili, Apache JMeter: A practical beginner's guide to automated testing and performance measurement for your websites, Packt Publishing Ltd, 2008.

[9] Apache JMeter, http://jmeter.apache.org/ (2014).

[10] Apache HTTP Server Benchmarking Tool, http://httpd.apache.org/docs/2.2/programs/ab.html (2014).

[11] HP LoadRunner, http://www8.hp.com/us/en/software-solutions/loadrunner-load-testing/ (2014).

[12] D. Draheim, J. Grundy, J. Hosking, C. Lutteroth, G. Weber, Realistic load testing of web applications, in: Software Maintenance and Reengineering, 2006. CSMR 2006. Proceedings of the 10th European Conference on, IEEE, 2006, pp. 11–pp.

[13] Custom Plugins for Apache JMeter, http://jmeter-plugins.org/ (2014).

[14] Y. Sun, J. White, S. Eade, A model-based system to automate cloud resource allocation and optimization, in: Model-Driven Engineering Languages and Systems, Springer, 2014, pp. 18–34.

[15] N. Newman, Apple ibeacon technology briefing, Journal of Direct, Data and Digital Marketing Practice 15 (3) (2014) 222–225.

[16] M. Kalin, Java web services: up and running, " O'Reilly Media, Inc.", 2013.

[17] Wrk Modern HTTP Benchmarking Tool, https://github.com/wg/wrk (2014).

[18] D. A. Bacigalupo, S. A. Jarvis, L. He, G. R. Nudd, An investigation into the application of different performance prediction techniques to e-commerce applications, in: Parallel and Distributed Processing Symposium, 2004. Proceedings. 18th International, IEEE, 2004, p. 248.

[19] AWS Cloud Formation, http://aws.amazon.com/cloudformation/ (2014).

[20] M. Jones, D. Hardt, The oauth 2.0 authorization framework: Bearer token usage, Tech. rep., RFC 6750, October (2012).

[21] Google Cloud Deployment Manager, https://cloud.google.com/deployment-manager/ (2014).

[22] Docker, https://www.docker.io (2014).

[23] Docker Hub Registry, https://registry.hub.docker.com/ (2014).

[24] M. Eysholdt, H. Behrens, Xtext: implement your language faster than the quick and dirty way, in: Proceedings of the ACM international conference companion on Object oriented programming systems languages and applications companion, ACM, 2010, pp. 307–309.

[25] L. Bettini, Implementing Domain-Specific Languages with Xtext and Xtend, Packt Publishing Ltd, 2013.

[26] J. Fink, Docker: a software as a service, operating system-level virtualization framework, Code4Lib Journal (25).

[27] AWS Elastic Beanstalk adds Docker support, http://aws.amazon.com/about-aws/whats-new/2014/04/23/aws-elastic-beanstalk-adds-docker-support/ (2014).

[28] Containers on Google Cloud Platform, https://cloud.google.com/compute/docs/containers (2014).

[29] W.-Y. Lin, G.-Y. Lin, H.-Y. Wei, Dynamic auction mechanism for cloud resource allocation, in: Cluster, Cloud and Grid Computing (CCGrid), 2010 10th IEEE/ACM International Conference on, IEEE, 2010, pp. 591–592.

[30] G. Wei, A. V. Vasilakos, Y. Zheng, N. Xiong, A game-theoretic method of fair resource allocation for cloud computing services, The Journal of Supercomputing 54 (2) (2010) 252–269.

[31] D. Warneke, O. Kao, Exploiting dynamic resource allocation for efficient parallel data processing in the cloud, Parallel and Distributed Systems, IEEE Transactions on 22 (6) (2011) 985–997.

[32] Elastic Load Balancing (ELB), https://aws.amazon.com/elasticloadbalancing/ (2014).

[33] N. Ferry, A. Rossini, F. Chauvel, B. Morin, A. Solberg, Towards model-driven provisioning, deployment, monitoring, and adaptation of multi-cloud systems, in: CLOUD 2013: IEEE 6th International Conference on Cloud Computing, 2013, pp. 887–894.

[34] L. Zhu, Y. Liu, N. B. Bui, I. Gorton, Revel8or: Model driven capacity planning tool suite, in: Software Engineering, 2007. ICSE 2007. 29th International Conference on, IEEE, 2007, pp. 797–800.

[35] T. Binz, U. Breitenbücher, O. Kopp, F. Leymann, Tosca: Portable automated deployment and management of cloud applications, in: Advanced Web Services, Springer, 2014, pp. 527–549.

[36] X. Wang, B. Zhou, W. Li, Model-based load testing of web applications, Journal of the Chinese Institute of Engineers 36 (1) (2013) 74–86.

[37] J. Li, J. Chinneck, M. Woodside, M. Litoiu, G. Iszlai, Performance model driven qos guarantees and optimization in clouds, in: Proceedings of the 2009 ICSE Workshop on Software Engineering Challenges of Cloud Computing, IEEE Computer Society, 2009, pp. 15–22.

[38] S. Chaisiri, B.-S. Lee, D. Niyato, Optimization of resource provisioning cost in cloud computing, Services Computing, IEEE Transactions on 5 (2) (2012) 164–177.

[39] S. Frey, F. Fittkau, W. Hasselbring, Search-based genetic optimization for deployment and reconfiguration of software in the cloud, in: Proceedings of the 2013 International Conference on Software Engineering, IEEE Press, 2013, pp. 512–521.

[40] M. Catan, R. Di Cosmo, A. Eiche, T. A. Lascu, M. Lienhardt, J. Mauro, R. Treinen, S. Zacchiroli, G. Zavattaro, J. Zwolakowski, Aeolus: Mastering the complexity of cloud application deployment, in: Service-Oriented and Cloud Computing, Springer, 2013, pp. 1–3.